

The price of anarchy and stability in general noisy best-response dynamics*

Paolo Penna¹

¹Department of Computer Science. ETH Zurich, Switzerland,
paolo.penna@inf.ethz.ch

December 15, 2015

Abstract

Logit-response dynamics [AFN10] are a rich and natural class of *noisy* best-response dynamics. In this work we revise the price of anarchy and the price of stability by considering the quality of long-run equilibria in these dynamics. Our results show that prior studies on simpler dynamics of this type can strongly depend on a *synchronous* schedule of the players' moves. In particular, a small noise by itself is not enough to improve the quality of equilibria as soon as other very natural schedules are used.

1 Introduction

Complex and distributed systems are often modeled by means of game dynamics in which agents (players) act spontaneously, and typically attempt to maximize their own benefit. These dynamics can be regarded as distributed algorithms that are implemented by the players and have applications in Economics (e.g., markets), Physics (e.g., Ising model and spin systems), Biology (e.g., evolution of life), and in Computer Science (e.g., distributed protocols for routing and resource allocations by competing entities). As a result of such selfish behavior the system will reach some equilibrium that may or may not be optimal. The PRICE OF ANARCHY (PoA) [KP09] and the PRICE OF STABILITY (PoS) [ADK⁺08] have been introduced to quantify the efficiency loss caused by the players selfish behavior. Intuitively, the price of anarchy compares the *worst* possible equilibrium to the optimum (the centrally designed outcome). Similarly, the price of stability compares the *best* possible equilibrium with the optimum (to capture the loss caused by the inherent instability of optimal outcomes). The PoA and the PoS measure the performance of “decentralized and anarchic” distributed systems in which we simply let the players decide by themselves how to play, until they reach some equilibrium.

*Part of this work has been done while at LIAFA, Université Paris Diderot, supported by the French ANR Project DISPLEXITY.

These two concepts seem naturally related to the question of explaining *how* the players can choose between different equilibria. A remarkable observation made in game theory is that *noise* allows the players to select good equilibria via a simple and natural dynamics. Also, results on the PoA and PoS for certain problems are sometimes based on equilibria that are “unnatural” because it is difficult for the players to reach or to remain in such equilibria. This leads naturally to reconsider the bounds on the PoA and PoS by restricting attention to “reasonable” equilibria.

In this work we follow this approach by considering the equilibria that arise in certain natural *noisy best-response* dynamics. In these dynamics, players keep updating their strategies and, at low noise regime, they tend to select strategies that deliver higher payoffs (see Section 1.3 for the formal model). The equilibria that are selected in the long-run are the so-called *stochastically stable* states. That is, the states that in the stationary distribution have positive probability as the noise vanishes.¹ One of the most popular and studied of such dynamics is the following one:

Logit dynamics [Blu93]. At each time step *one randomly chosen player* updates her strategy according to a noisy best-response rule.

On the one hand, for potential games, the long-run outcomes can be simply characterized as the Nash equilibria *minimizing the potential* of the game [Blu98]. On the other hand these dynamics seem to require a particular *schedule* of the players’ moves (who is allowed to update at each step and who stays put). The authors in [AFN10] analyze the generalization in which players’ moves are not necessarily “synchronized”:

Independent learning logit-response dynamics [AFN10]. At each time step a *subset* of randomly chosen players update their strategies according to a noisy best response-rule (the same rule of logit dynamics). Every subset has some positive probability of being selected (e.g., when every player independently decides to revise his/her strategy with some probability $p \in (0, 1)$).

While the characterization of [Blu98] in terms of potential minimizers is no longer true, [AFN10] characterize long-run outcomes as the states minimizing a *stochastic potential*. This means that the set of equilibria selected by the two dynamics is *different* [AFN10]. It is therefore natural to ask to what extent the *schedule* of the players’ moves (the *revision process*, in the game theory terminology) can affect the *quality* of the selected equilibria in these dynamics. The logit-response dynamics has been proven to maintain many of the nice properties of logit dynamics, including the selection of good equilibria in coordination games, or convergence to Nash equilibria under certain conditions [AFN10, MS12, CDGT14, AFN15].

1.1 Our Contribution

This work presents a refined version of the quality of equilibria in games (PoA and PoS) based on ‘decentralized’ dynamics with noise, that is, the *logit-response with independent learning* described above. This approach has two advantages:

¹In this work we consider *pure* Nash equilibria, that is, states in which each player chooses one strategy and unilateral deviations are not beneficial. In general games, stochastically stable states are not necessarily Nash equilibria, even for the class of potential games (which is the main focus of this work).

1. It suggests an intuitive definition of ‘reasonable’ equilibria (robust to small mistakes and easy to reach without specific schedules of players’ moves) to quantify the efficiency loss.
2. It allows to understand the role of *synchronization* in such dynamics (prior studies considered the quality of equilibria in ‘synchronous’ logit dynamics).

Specifically, we propose to study a natural concept of INDEPENDENT-LOGIT POA (and POS) which compares the worst (respectively, best) long-run outcomes of independent learning logit-response dynamics with the optimum. That is, the classical POA (respectively, POS) restricted to the set of stochastically stable states [AFN10]. The analogous notions for logit dynamics, say the LOGIT POA and LOGIT POS, have been already studied in the literature under different names.² A direct comparison between these notions is informative of the importance of the schedule of the moves.

We study two classes of problems that have different features: *load balancing* games which provide a simple model of congestion, and *broadcast network design* games in which players share the cost of the used resources. Intuitively, in the first case players dislike to choose the same resources (because this creates congestion and players experience some delay), while in the second case they like to do the opposite (because the more players use a resource the less each of them pays for it).

We first consider *load balancing* games whose analysis for logit dynamics has been done by [AS09]. We obtain the tight bounds shown in Table 1 which illustrate of how different the two dynamics can perform. In particular, though the prior bounds on logit dynamics show that the quality of equilibria *improves* compared to “generic” Nash equilibria, the opposite happens for independent-learning dynamics. This tells us that the improvement of the quality of the equilibria is due to the *combination* of the players’ response rule and of the possibility of *synchronizing* their moves. Thus, the bounds based on potential minimizers are in general not “robust” in the sense that they depend critically on this assumption. Also, contrary to logit dynamics, the stochastically stable states are not necessarily Nash equilibria (in fact they are far from Nash equilibria in terms of cost).

We then consider *broadcast network design* games studied in [KM13, MM15] for potential minimizers (logit dynamics), and show that in this case synchronization is not essential to obtain optimal configurations. First, for the case of parallel links between a common source and a common destination (links represent resources of different costs) players will select the optimum if there is a unique shortest link (Theorem 3.2). Note that in these instances the price of anarchy is rather high ($\text{POA} = n$), while our result says that the restriction to stochastically stable states is instead optimal ($\text{INDEPENDENT-LOGIT POA} = 1$). Finally, we show that the result does not extend to different sources since in some instances $\text{INDEPENDENT-LOGIT POA} > \text{LOGIT POA}$ (Theorem 3.1). Also in this case stochastically stable states are not Nash equilibria.

²[AS09] call the POA restricted to potential minimizers INEFFICIENCY RATIO OF STABLE EQUILIBRIA (IRSE), while [KM13] use the terms POTENTIAL OPTIMAL PRICE OF ANARCHY (POPoA) and OF STABILITY (POPoS). In this work we prefer to use the terms LOGIT POA and LOGIT POS to emphasize the comparison between logit and independent-learning logit-response dynamics.

Noise with Synchronization	Deterministic	Noise without Synchronization
Worst Equilibrium		
LOGIT PoA [AS09]	PoA [FH79]	INDEPENDENT-LOGIT PoA
$\geq 19/18$	$2\left(1 - \frac{1}{m+1}\right)$	m
$\leq 3/2$		
Best Equilibrium		
LOGIT PoS	PoS	INDEPENDENT-LOGIT PoS
1	1	$2\left(1 - \frac{1}{m+1}\right)$

Table 1: Bounds for load balancing games: The rightmost column contains our contributions, and m denotes the number of machines. The bounds on (LOGIT) PoS are folklore.

1.2 Further Related Work

The impact of the schedule of the players’ moves on the quality of the equilibria has been studied earlier. The approach by [MS12] is closely related to the one here, and it consists of deriving sufficient conditions for which the stochastically stable states of logit-response dynamics are the same as those of logit dynamics (potential minimizers). The impact of simultaneous moves on (deterministic) best-response dynamics is studied by [FMS12]. The performance of various *concurrent* dynamics in which all players are active at each time step is studied in e.g. [FKS10, ABFH09, KPT09, AFP+13].

Several works refine the PoA and the PoS by restricting to certain type of equilibria. The *price of stochastic anarchy* by [CLPR08, CP09] is similar to ours in the sense that it considers stochastically stable states though for a (concurrent) imitation dynamics. The authors in [AS09] were the first to restrict to potential minimizers (the long-run equilibria of logit dynamics) and their analysis on atomic congestion games and load balancing games shows that prior bounds on the PoA for these games can be too pessimistic. The work [KM13] extends this approach to the PoS and shows similar results for broadcast network design games (where LOGIT PoS = LOGIT PoA up to a constant factor). The restriction to the ring topology is considered in [MM15] (together with other variants of PoS and PoA). The authors in [AFPP13] suggest to consider the expected social cost in logit dynamics and therefore the stationary distribution as the equilibrium concept. A new natural dynamics converging to Nash equilibria with optimal social welfare in general games has been proposed in [PY12]. The robustness of the PoA and PoS bounds has been investigated by [CKS11] and [Rou09] who considered approximate and correlated equilibria, respectively.

Stochastic stability is the tool to show that in the coordination game players select the risk dominant strategy (see [AFN10] for discussion and references). Recent works have pointed out that in general revision processes (including independent-learning) logit-response dynamics do *not* converge to Nash equilibria in potential games ([CDGT14, AFN15] discuss this issue explicitly).

1.3 Preliminary Definitions

This section introduces the necessary tools from [AFN10], while the next subsection presents the revised PoA and PoS notions. Let $u_i(s_i, s_{-i})$ denote the utility of player i when she chooses strategy s_i and the other players' strategies are s_{-i} (the latter is a vector of strategies). The combination of all these strategies is a state $s = (s_i, s_{-i})$. In logit-response dynamics, whenever a player is given the opportunity to revise her strategy, she will revise her current strategy with a strategy s_i chosen according to the *logit choice function*

$$p_i(s_i, s_{-i}) = \frac{e^{\beta u_i(s_i, s_{-i})}}{\sum_{s'_i \in S_i} e^{\beta u_i(s'_i, s_{-i})}} \quad (1)$$

where s_{-i} is the current strategy profile (the strategies chosen by the others) and β is a parameter representing the inverse level of noise in the players' decision ($0 < \beta < \infty$), and S_i is the set of strategies of player i . The dynamics is defined as follows:

- Select a subset J of players according to a probability distribution q defined on the subsets of players.
- Every selected player $i \in J$ chooses a new strategy s_i with probability given by the logit choice function (1).

Note that the dynamics depends on the *revision process* q , and different revision processes yields different (logit-response) dynamics:

- Logit dynamics [Blu93, Blu98] correspond to *asynchronous learning* in which only one player at the time is allowed to revise (i.e, $q(J) > 0$ if and only if $J = \{i\}$ for some i).
- In logit-response dynamics with *independent learning* every subset of players has non-zero revising probability (i.e. $q(J) > 0$ for every subset J).

A state s is *stochastically stable* if it has positive probability in the limit invariant distribution as the noise goes to zero,

$$\lim_{\beta \rightarrow \infty} \mu^\beta(s) > 0,$$

where μ^β is the stationary distribution of the dynamics with parameter β .³

An edge (s, s') is *feasible* if $R_{s,s'} := \{J \mid q(J) > 0 \text{ and } s_k = s'_k \forall k \notin J\}$ is non-empty (note that $R_{s,s'}$ consists of the subsets of players potentially leading from s to s'). For any such feasible transition, its *waste* is defined as

$$W_{s,s'} := \min_{J \in R_{s,s'}} W_{s,s'}^{(J)}, \text{ where}$$

$$W_{s,s'}^{(J)} := \sum_{j \in J} \left(\max_{s''_j \in S_j} u_j(s''_j, s_{-j}) - u_j(s'_j, s_{-j}) \right).$$

³It is well known that the dynamics described above define an ergodic Markov chain over the possible strategy profiles for every $0 < \beta < \infty$. That is, the stationary distribution μ^β exists for all β .

The waste of a directed tree (or of path) T is simply the sum of the waste of its edges:

$$W(T) := \sum_{(s,s') \in T} W_{s,s'}.$$

The *stochastic potential* of a state s is

$$W(s) := \min_{T \in \mathcal{T}(s)} W(T)$$

where $\mathcal{T}(s)$ is the set of all directed trees T in which every state s'' has a unique path to s . That is, T is a tree directed towards s with each edge being feasible.

Theorem 1.1 ([AFN10]). *Consider the logit-response dynamics (with any revision process). A state is stochastically stable if and only if it minimizes $W(s)$ among all states.*

It has been observed recently that, even for potential games⁴, stochastically stable states are not necessarily Nash equilibria [CDGT14, AFN15].

In our proofs we shall find it convenient to reason about the existence of *zero-waste* paths between two states. Note that a zero-waste transition (s, s') corresponds to a multiple best-response: Since $W_{s,s'}^{(J)} = 0$ all players in J simultaneously best-respond to the strategies in s and the resulting profile is s' (for $J = \{i\}$ we have a best-response).

1.3.1 Revised Price of Anarchy and Stability

We want to quantify the efficiency loss that occurs when we relax logit dynamics to independent learning revision processes. Note that the set of stochastically stable states does not depend on the particular revision process q satisfying the definition of independent learning (the stochastic potential depends only on the feasible transitions which do not change if q and q' are both independent learning revision processes). Given an instance I (a game) of interests, and a corresponding global cost function *cost* mapping states into a nonnegative real, we define the following

$$\text{INDEPENDENT-LOGIT POA}(I) = \frac{\max_{s \in \text{STABLE}} \text{cost}(s, I)}{\min_s \text{cost}(s, I)}$$

where *STABLE* denotes the stochastically stable states for the logit-response dynamics with independent learning. By replacing ‘max’ with ‘min’ we obtain the notion of INDEPENDENT-LOGIT POS. The original POA and POS are defined by replacing *STABLE* with the set of all possible Nash equilibria. The restriction to logit dynamics, LOGIT-POA and LOGIT-POS, have *STABLE* equal to the states minimizing the potential of the game. We generally consider classes of games and extend both definitions by taking the supremum over all instances: For a class of games \mathcal{C} the INDEPENDENT-LOGIT POA is defined as $\sup_{I \in \mathcal{C}} \text{INDEPENDENT-LOGIT POA}(I)$; INDEPENDENT-LOGIT POA is defined as $\sup_{I \in \mathcal{C}} \text{INDEPENDENT-LOGIT POS}(I)$.

The following easy result (whose elementary proof is given in Appendix A.1) is useful for games in which convergence to Nash equilibria is not guaranteed (like those considered here).

⁴In this work we deal only with weighted potential games, that is, games that admit a vector w and a potential function ϕ such that $u_i(s) - u_i(s') = (\phi(s') - \phi(s))w_i$ for all i and for all s and s' that differ in exactly player i ’s strategy.

Theorem 1.2. *For any game, the set of stochastically stable states in independent learning logit-response dynamics must contain at least one Nash equilibrium. Therefore*

$$\text{INDEPENDENT-LOGIT PoS} \leq \text{PoA}.$$

2 Load Balancing Games

We consider the load balancing problem which is defined as follows. There are m identical machines and n jobs, each job having its own weight. Each job is controlled by a player who can choose one machine (strategy). The strategies determine an allocation and the load for each machine (sum of the weights of jobs allocated to the machine). Naturally, each player aims at having her job on a machine with minimal load (so the payoff is defined as the negative of the load of the chosen machine). The cost of an allocation is the so-called *makespan*, that is, the maximum load over all machines.

2.1 Independent Logit PoA

In this section we show that the PoA in independent learning dynamics can be much worse than all prior bounds. In particular, it is higher than the classical PoA which considers any Nash equilibria. This means that there are some stochastically stable states which are *not* Nash equilibria for this potential game.

Theorem 2.1. *For load balancing games,*

$$\text{INDEPENDENT-LOGIT PoA} = m.$$

Proof. Since the upper bound is trivial, we shall prove only the lower bound. Consider the instance with m machines and $n = lm - 1$ identical jobs (say of size 1), for a large positive integer l . The optimum has cost l and it is given by any allocation in which one machine has $l - 1$ jobs, and all other machines have l jobs.

We claim that *every* state is stochastically stable, including those in which *all* jobs go to the same machine (and whose cost is thus $lm - 1$). To prove our claim, we observe that in every state s there is at least one machine m_{high} with at least l jobs and another machine m_{low} with at most $l - 1$ jobs. This implies that we can construct a zero-waste path from any state s to every other state s' as follows:

1. First all jobs in machine m_{high} move to machine m_{low} ;
2. Now that machine m_{high} is empty, all jobs move there;
3. Then from there all jobs move to the machine they occupy in s' (some do not move because they are already in the correct machine).

Note that each move is a best-response made simultaneously by a subset J of players (because players move to a machine with strictly lower load). Therefore each of the three transitions has zero waste. Since the stochastically stable state with all jobs in the same machine costs $lm - 1$, and the optimum is l , we have shown that

$$\text{INDEPENDENT-LOGIT PoA} \geq m - 1/l$$

for all l . The theorem thus follows by taking l large enough. □

2.2 Independent Logit PoS

In this section we show a tight bound on the INDEPENDENT-LOGIT PoS.

The strategy to prove that a certain state is *not* stochastically stable is to prove that it is “easy” to leave this state. That is, there exists a tree of small waste which is directed *from* this state to all other states. The next lemma says that such a tree can be obtained by finding a small-waste path leading to a state in which one machine has *zero load*.

Lemma 2.2. *Let s be an allocation in which one machine has zero load. Then there exists a zero-waste path from s to every other state.*

The proof of this lemma is exactly the same argument used in the proof of Theorem 2.1 above (Items 2-3). Next, we introduce some additional notation that will ease the proofs below.

Compact notation for states and jobs allocations. Given a set of jobs weights, we partition these weights among the machines, without distinguishing among the machines and among players having identical weights. This results in a *class* of states having the same cost. For instance, for jobs of size $\{2, 1, 1, 1\}$ the partition

$$[2], [1, 1], [1]$$

denotes a *class* of allocations in which one machine gets a job of size 2, does not matter which of the three machines, another machine gets two out of the three jobs of size 1, and the remaining machine gets the remaining job of size 1. An explicit representation would be in terms of which player goes to which machine: for players $\{p_1, p_2, p_3, p_4\} = \{2, 1, 1, 1\}$ the states

$$[p_1], [p_2, p_3], [p_4], \quad [p_1], [p_2, p_4], [p_3], \quad [p_1], [p_3, p_4], [p_1], \dots, [p_4], [p_2, p_3], [p_1]$$

correspond to the partition $[2], [1, 1], [1]$. Finally, we call a job of size s an s -job. \square

Theorem 2.3. *For load balancing games,*

$$\text{INDEPENDENT-LOGIT PoS} \geq 2 \left(1 - \frac{1}{m+1} \right).$$

Proof. Consider instances with m machines and jobs

$$\{\Delta - \delta, \Delta - \delta, \underbrace{\Delta, \dots, \Delta}_{m-2}, \underbrace{\delta, \dots, \delta}_{lm}\}$$

where l is a (sufficiently large) positive integer and

$$\Delta = \frac{m}{m+1}, \quad \text{and} \quad \delta = \frac{\Delta}{lm}.$$

First, optimal allocations have cost 1 and can be obtained (only) by allocating one large job (size Δ or $\Delta - \delta$) in each machine as follows:

$$[\Delta - \delta, \underbrace{\delta, \dots, \delta}_l], [\Delta - \delta, \underbrace{\delta, \dots, \delta}_l], [\Delta, \underbrace{\delta, \dots, \delta}_l], \dots, [\Delta, \underbrace{\delta, \dots, \delta}_l]. \quad (2)$$

(Note that $\Delta + l\delta = 1$ and the cost is indeed 1.)

Second, there are Nash equilibria of cost $2(\Delta - \delta)$ which are of the form

$$[\Delta - \delta, \Delta - \delta], [\delta, \dots, \delta], [\Delta], \dots, [\Delta]. \quad (3)$$

We shall prove that only allocations of cost $2(\Delta - \delta)$ are stochastically stable, thus implying

$$\text{INDEPENDENT-LOGIT POS} \geq 2(\Delta - \delta) = 2\left(1 - \frac{1}{m+1} - \delta\right).$$

The theorem thus follows by taking δ arbitrarily small, i.e., l arbitrarily large.

Claim 2.4. *Every allocation with one large job per machine has a zero-waste path to every other state.*

Proof of Claim. In any allocation that allocates one large job per machine, there exists a machine with load at most $\Delta - \delta + l\delta = 1 - \delta$. The following sequence of best-response moves leads to another allocation with one machine having zero load:

$$\begin{array}{ccccccc} [\Delta - \delta, \underbrace{\delta, \dots, \delta}_l, & [\Delta - \delta, \underbrace{\delta, \dots, \delta}_l, & [\Delta, \underbrace{\delta, \dots, \delta}_l, \dots, & [\Delta, \underbrace{\delta, \dots, \delta}_l & \rightarrow \\ [\Delta - \delta, \underbrace{\delta, \dots, \delta}_l, & [\Delta - \delta, \underbrace{\delta, \dots, \delta}_{l(m-1)}, & [\Delta], \dots, & [\Delta] & \rightarrow \\ [\Delta - \delta], & [\Delta - \delta], & [\Delta], \dots, & [\Delta, \underbrace{\delta, \dots, \delta}_{lm} & \rightarrow \\ [\Delta - \delta, \Delta, \underbrace{\delta, \dots, \delta}_{lm}, & [\Delta - \delta], & [\Delta], \dots, & [\text{empty}] \end{array}$$

The claim then follows by Lemma 2.2. \square

Claim 2.5. *Every Nash equilibrium of cost $2(\Delta - \delta)$ requires a minimal waste of δ and there exists a path of waste δ from such an equilibrium to every other state.*

Proof of Claim. Every Nash equilibrium of cost $2(\Delta - \delta)$ must allocate two $(\Delta - \delta)$ -jobs to the same machine and thus it must be of the form:

$$[\Delta - \delta, \Delta - \delta], [\delta, \dots, \delta], [\Delta], \dots, [\Delta].$$

Since the machine containing all δ -jobs has load Δ , every move of one (or more) job(s) requires a waste of δ or larger. To get a path of waste δ , consider the following transitions:

$$\begin{array}{ccccc} [\Delta - \delta, \Delta - \delta], & [\delta, \dots, \delta], & [\Delta], \dots, [\Delta] & \xrightarrow{(waste=\delta)} \\ [\Delta - \delta], & [\Delta - \delta, \delta, \dots, \delta], & [\Delta], \dots, [\Delta] & \xrightarrow{(waste=0)} \\ [\Delta - \delta, \Delta - \delta, \delta, \dots, \delta], & [\text{empty}], & [\Delta], \dots, [\Delta] \end{array}$$

and apply Lemma 2.2 to obtain the rest of the path. \square

To prove the theorem, let APX denote the approximate allocations of the form (3). We use the two claims above to prove that for every $s \notin APX$

$$W(s) \geq |APX|\delta \quad (4)$$

while for every $s_{apx} \in APX$

$$W(s_{apx}) \leq (|APX|-1)\delta \quad (5)$$

Inequality (4) follows from the fact that every state in APX requires a waste of δ to leave, thus any tree towards s must have waste $|APX|\delta$. For (5), we argue as follows. First we build a partial tree from the $|APX|-1$ states in APX to s_{apx} . Second, we observe that every non-Nash equilibrium has a zero-waste path to some Nash equilibrium. Third, we show that set of all Nash equilibria is the union of APX and OPT , where OPT are the allocations of cost 1 (details in Appendix A.2). Then, by Claim 2.4, the paths from OPT to s_{apx} have zero waste. \square

A matching upper bound follows from the known upper bound on PoA (see Table 1) and the fact that INDEPENDENT-LOGIT PoS \leq PoA (Theorem 1.2).

Corollary 2.6. *For load balancing games,*

$$\text{INDEPENDENT-LOGIT PoS} = 2\left(1 - \frac{1}{m+1}\right).$$

3 Broadcast Network Design Games

In this section we consider broadcast network design games studied in [KM13, MM15]. In these games, we are given a network in which edges have costs, and a set of players located in some of the nodes. The players aim at connecting to a common terminal node t and their strategy is to choose which route to use. The cost of each edge is then *shared equally* among all players using that edge: if n_e players select a route containing edge e then each of them pays a share c_e/n_e of the cost c_e of the edge. The cost for a player is the sum of the shares of the edges contained in the chosen route, and the social cost is the sum of the costs of the used edges.

Figure 1 shows an example from [KM13] in which s_1 and s_2 denote the two players (their locations). We use this instance to show that also in these games synchronization is important for the quality of equilibria.

Theorem 3.1. *There exists an instance I of a broadcast network design game for which $\text{PoA}(I) = 4/3$ but*

$$\text{INDEPENDENT-LOGIT PoA}(I) = 5/3.$$

In particular, some non-Nash equilibrium is stochastically stable.

Proof. Consider the instance in Figure 1. Each player has two strategies and therefore there are only four possible states shown in Figure 1. Note that the Nash equilibria are s^1, s^2 and s^3 , and s^1 and s^3 can be obtained from s^2 by a move of a single player whose cost remains 2 in both cases (thus the move is a best-response). Therefore, transitions




Figure 1: On the left, an instance of broadcast network design game for which a non-Nash equilibrium is also stochastically stable. On the right, the possible states for this instance (think arrows represent zero-waste transitions).

(s^2, s^1) and (s^2, s^3) are zero-waste transitions. Moreover, starting from s^2 , if both players move we obtain state s^0 and again a zero-waste transition (s^2, s^0) (being each move a best-response). Finally, transitions from s^0 to any Nash equilibrium are zero-waste because the player(s) changing strategy actually improve. Since all transitions are zero-waste, we have $W(s) = 0$ for all four states and therefore s^0 is stochastically stable. Since the cost of s^0 is 5 and the optimum is 3, the theorem follows. \square

We next consider broadcast network design games *over parallel links*: All players are located on *the same node* s and the graph consists of m parallel links (representing resources of different costs) connecting to the common target t . Let $\ell_1 \leq \ell_2 \leq \dots \leq \ell_m$ be the costs of the links, and let N_k denote the state in which all players choose the k^{th} link in the order above (in particular N_1 is optimal and also a Nash equilibrium). The following result provides sufficient conditions for which the dynamics will select only the optimum.

Theorem 3.2. *For network broadcast games over parallel links, if there is a unique shortest link ($\ell_1 < \ell_2$) or all links have the same cost then*

$$\text{INDEPENDENT-LOGIT POA} = 1.$$

The proof uses the following *radius-coradius* argument from [AFN10]. The *basin of attraction* of a state s is the set $B(s)$ of all states s' such that there is a zero-waste path from s to s' . The *radius* $R(s)$ is the minimum number r such that there is a path of waste r from s to some $s' \notin B(s)$. The *coradius* $CR(s)$ is the minimum number cr such that, for every $s'' \notin B(s)$, there is a path of waste at most cr from s'' to s .

Lemma 3.3 ([AFN10]). *If $R(s) > CR(s)$, then the set of stochastically stable states are exactly those in $L(s)$, where $L(s)$ is the set of $s' \in B(s)$ for which $s \in B(s')$.*

We next apply this result to our problem.

Lemma 3.4. $R(N_1) - CR(N_1) = (\ell_2 - \ell_1)H(n)$, where $H(n) = 1 + 1/2 + \dots + 1/n$ is the harmonic function.

Proof. Let b_1 be the smallest positive integer such that if link 1 is chosen by at least b_1 players, then moving to link 1 is a best-response for all other players. That is

$$\frac{\ell_1}{b_1 + 1} \leq \frac{\ell_2}{n - b_1}.$$

Claim 3.5. $R(N_1) \geq (\ell_2 - \frac{\ell_1}{n}) + (\frac{\ell_2}{2} - \frac{\ell_1}{n-1}) + \dots + (\frac{\ell_2}{n-b_1} - \frac{\ell_1}{b_1-1})$.

Proof of Claim. By definition of b_1 , in order to reach a state not in $B(N_1)$ we have to move $n - b_1 + 1$ players from link 1. The minimal waste path for achieving this is to move these players one by one and all to link 2. Note that the k^{th} move has waste $\frac{\ell_2}{k} - \frac{\ell_1}{n-k+2}$ and the last move ($k = n - b_1 + 1$) has zero waste by definition of b_1 : indeed $\frac{\ell_1}{b_1} > \frac{\ell_2}{n-b_1+1}$. \square

Claim 3.6. $CR(N_1) \leq (\ell_1 - \frac{\ell_2}{n}) + (\frac{\ell_1}{2} - \frac{\ell_2}{n-1}) + \dots + (\frac{\ell_1}{b_1} - \frac{\ell_2}{n-b_1+1})$

Proof of Claim. By definition of b_1 , starting from any state, it is enough to move at most b_1 players into link 1 to obtain a state is in $B(N_1)$. We move players one by one. Suppose first that initially ℓ_1 contains no players. Then the k^{th} move has waste $\frac{\ell_1}{k} - \frac{\ell_{j_k}}{n_{j_k}}$, where $\ell_{j_k} \neq \ell_1$ is the link of the k^{th} player we move, and n_{j_k} is the number of players currently in that link. Since $\ell_{j_k} \geq \ell_2$ and $n_{j_k} \leq n - k + 1$ (link 1 contains $k - 1$ before this move is done), the waste of each move is upper bounded by $\frac{\ell_1}{k} - \frac{\ell_2}{n-k+1}$ and a total of b_1 moves is enough. This gives the upper bound in the claim. Finally, if ℓ_1 already contained some players in the initial state, then an even smaller waste is sufficient (it suffice to consider the first n_1 moves above as already done, where n_1 is the initial number of players in ℓ_1). \square

By combining the two claims above and rearranging terms, we get $R(N_1) - CR(N_1) \geq \ell_2 + \frac{\ell_2}{2} + \dots + \frac{\ell_2}{n} - (\ell_1 + \frac{\ell_1}{2} + \dots + \frac{\ell_1}{n})$, which proves the lemma. \square

Proof of Theorem 3.2. If all links have the same cost ($\ell_1 = \ell_m$) then every state N_k is optimal. Since these states are also strict Nash equilibrium, the convergence result in [AFN15] says that the stochastically stable states are contained in these strict Nash equilibria (formally, the game is weakly acyclic since every other state is not a Nash equilibrium).

Now consider the case in which there is only *one* optimal link ($\ell_1 < \ell_2$), then we have $R(N_1) > CR(N_1)$ by Lemma 3.4. Moreover $L(N_1) = \{N_1\}$ since this is a strict Nash equilibrium. We can thus apply Lemma 3.3. \square

4 Conclusion and Open Questions

In this work we have studied a refinement of the POA and POS based on the long-run equilibria in a general noisy best-response dynamics (a generalization of logit dynamics). The analysis is based on the characterization by [AFN10] and it provides new insights on the importance of the order in which players revise their strategies. On the one hand, restricting the analysis only to potential minimizers may not be significant for other natural dynamics of the same type. On the other hand, a comparison between the two dynamics is useful to understand the importance of synchronous moves.

We remark that, in load balancing games, the bounds for independent learning logit-response dynamics are in a sense the *worst possible*. First, m is an upper bound to *any* configuration, and therefore INDEPENDENT-LOGIT POA $\leq m$. Second, in every game we have INDEPENDENT-LOGIT POS \leq POA (see Theorem 1.2). In load balancing games both inequalities are tight.

A natural alternative would be to consider the expected cost at stationary distribution [AFPP13]. Here the main technical difficulty lies in obtaining the stationary distribution (see e.g. [AFP⁺13] for such an issue). The advantage of stochastic stability is that it avoids an explicit computation of the stationary distribution. However, the resulting bounds can be considered too pessimistic in some cases. For instance, when each player is selected for revising her strategy with probability p sufficiently small, the dynamics should be close to the logit dynamics. [MS12] suggest to reduce to potential minimizers, that is, to show that stochastically stable states are potential minimizers also for logit-response dynamics with more general revision processes. We note that this approach requires additional hypothesis on the game (or on the players behavior). In particular, one should assume that players have a unique best response to have convergence to Nash equilibria [CDGT14, AFN15]. In this case, one gets INDEPENDENT-LOGIT PoA \leq PoA, and in some cases the results of [MS12] say that INDEPENDENT-LOGIT PoA \leq LOGIT PoA.

The study of other problems considered in [AS09, KM13, MM15] is an interesting direction for future research. Some of these bounds on potential minimizers are already quite sophisticated, and stochastic stability in independent learning may require non-trivial analysis. Estimating INDEPENDENT-LOGIT PoA and POS is still open even for ring networks [MM15].

Acknowledgments. I am grateful to Francesco Pasquale for comments on an earlier version of this work, and to Carlos Alos-Ferrer for bringing [CDGT14, AFN15] to my attention.

A Postponed Proofs and Additional Details

A.1 Proof of Theorem 1.2

Observe that for any state \bar{s} which is not a Nash equilibrium there exists a sequence of best-response moves that reaches some Nash equilibrium s . That is, there exists a zero-waste path from any non-Nash \bar{s} to some Nash s .

This implies that $W(\bar{s}) \geq W(s)$ because any waste tree directed to \bar{s} can be transformed into a tree of no larger waste directed into s (by adding the path mentioned above and removing other transitions). This shows that the waste $W(\bar{s})$ cannot be smaller than the waste of all Nash equilibria.

A.2 Postponed Details for the Proof of Theorem 2.3

We need to prove that the set of all Nash equilibria is just the union of *OPT* and *APX*, that is, the allocations of the form (2) and (3), respectively.

We first show that any allocation with a machine containing a Δ -job and another large job (size Δ or $\Delta - \delta$) cannot be a Nash equilibrium. To be a Nash equilibrium, the load of every other machine should be at least Δ . This will not be possible given the remaining jobs: If we have a machine with a Δ -job and $\Delta - \delta$ -job, then the jobs that are

left for the other $m - 1$ machines are

$$\{\Delta - \delta, \underbrace{\Delta, \dots, \Delta}_{m-3}, \underbrace{\delta, \dots, \delta}_{lm}\}$$

where the δ -jobs sum up to Δ . Thus the total load we can distribute is $(m - 1)\Delta - \delta$, and some machine will have load less than Δ . The same holds if we start with a machine having two Δ -jobs.

Finally, we show that every allocation with one large job per machine has a zero-waste path to some allocation in *OPT*. Indeed, allocations with one large job per machine are of the form

$$[\Delta - \delta, \underbrace{\delta, \dots, \delta}_{k_1}], [\Delta - \delta, \underbrace{\delta, \dots, \delta}_{k_2}], [\Delta, \underbrace{\delta, \dots, \delta}_{k_3}], \dots, [\Delta, \underbrace{\delta, \dots, \delta}_{k_m}]. \quad (6)$$

To obtain an allocation in *OPT* it is enough to move the δ -jobs from higher load machines to lower load machines (or to machines with the same load), until we obtain $k_1 = \dots = k_m = l$, that is (2).

References

- [ABFH09] Heiner Ackermann, Petra Berenbrink, Simon Fischer, and Martin Hoefer. Concurrent imitation dynamics in congestion games. In *Proc. of the 28th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 63–72, 2009.
- [ADK⁺08] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.
- [AFN10] Carlos Alós-Ferrer and Nick Netzer. The logit-response dynamics. *Games and Economic Behavior*, 68(2):413–427, 2010.
- [AFN15] Carlos Alós-Ferrer and Nick Netzer. On the convergence of logit-response to (strict) nash equilibria. TR available at Research Gate, June 2015.
- [AFP⁺13] Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, Paolo Penna, and Giuseppe Persiano. Logit dynamics with concurrent updates for local interaction games. In *Proc. of the 21st Annual European Symp. on Algorithms (ESA)*, volume 8125 of *LNCS*, pages 73–84, 2013.
- [AFPP13] Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, and Giuseppe Persiano. Mixing time and stationary expected social welfare of logit dynamics. *Theory of Computing Systems*, 53(1):3–40, 2013.
- [AS09] Arash Asadpour and Amin Saberi. On the inefficiency ratio of stable equilibria in congestion games. In *Proc. of the 5th Int. Workshop on Internet and Network Economics (WINE)*, volume 5929 of *LNCS*, pages 545–552, 2009.

- [Blu93] Lawrence E. Blume. The statistical mechanics of strategic interaction. *Games and economic behavior*, 5(3):387–424, 1993.
- [Blu98] Lawrence E. Blume. *Population games*. Addison-Wesley, 1998.
- [CDGT14] Pierre Coucheney, Stéphane Durand, Bruno Gaujal, and Corinne Touati. General revision protocols in best response algorithms for potential games. In *Network Games, Control and Optimization (NetGCoop)*, 2014.
- [CKS11] George Christodoulou, Elias Koutsoupas, and Paul G Spirakis. On the performance of approximate equilibria in congestion games. *Algorithmica*, 61(1):116–140, 2011.
- [CLPR08] Christine Chung, Katrina Ligett, Kirk Pruhs, and Aaron Roth. The price of stochastic anarchy. In *Proc. of the 1st Int. Symp. on Algorithmic Game Theory (SAGT)*, volume 4997 of *LNCS*, pages 303–314, 2008.
- [CP09] Christine Chung and Evangelia Pyrga. Stochastic stability in internet router congestion games. In *Proc. of the 2nd Int. Symp. on Algorithmic Game Theory (SAGT)*, volume 5814 of *LNCS*, pages 183–195, 2009.
- [FH79] Greg Finn and Ellis Horowitz. A linear time approximation algorithm for multiprocessor scheduling. *BIT Numerical Mathematics*, 19(3):312–320, 1979.
- [FKS10] Dimitris Fotakis, Alexis C Kaporis, and Paul G Spirakis. Atomic congestion games: Fast, myopic and concurrent. *Theory of Computing Systems*, 47(1):38–59, 2010.
- [FMS12] Angelo Fanelli, Luca Moscardelli, and Alexander Skopalik. On the impact of fair best response dynamics. In *Proc. of the 37th international conference on Mathematical Foundations of Computer Science (MFCS)*, volume 7464 of *LNCS*, pages 360–371, 2012.
- [KM13] Yasushi Kawase and Kazuhisa Makino. Nash equilibria with minimum potential in undirected broadcast games. *Theoretical Computer Science*, 482:33–47, 2013.
- [KP09] Elias Koutsoupas and Christos H. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009.
- [KPT09] Robert Kleinberg, Georgios Piliouras, and Éva Tardos. Multiplicative updates outperform generic no-regret learning in congestion games. In *Proc. of the 41st Annual ACM Symp. on Theory of Computing (STOC)*, pages 533–542, 2009.
- [MM15] Akaki Mamageishvili and Matús Mihalák. Multicast network design game on a ring. *CoRR*, abs/1507.04222, 2015. To appear in COCOA 2015.
- [MS12] Jason R. Marden and Jeff S. Shamma. Revisiting log-linear learning: Asynchrony, completeness and payoff-based implementation. *Games and Economic Behavior*, 75(2):788–808, 2012.

- [PY12] Bary S.R. Pradelski and H. Peyton Young. Learning efficient nash equilibria in distributed systems. *Games and Economic behavior*, 75(2):882–897, 2012.
- [Rou09] Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *Proc. of the 41st ACM Symp. on Theory of Computing (STOC)*, pages 513–522, 2009.